Decentralized Decision Making in Mean Field Stochastic Systems

Minyi Huang

School of Mathematics and Statistics
Carleton University
Ottawa, Canada

EPSRC Workshop: Topics in Control
University of Warwick, Nov. 30 – Dec. 2, 2011
Joint work with P.E. Caines, R.P. Malhamé, and S.L. Nguyen

Backgrounds of mean field decision making

- Production planning (Lambson)
- Advertising competition (Erickson)
- Wireless network resource management (HCM; Tembine et. al.)
- Industry dynamics with many firms (Weintraub, Benkard, Roy; Adlakha, Johari, Goldsmith)
- Selfish herding (such as fish) reducing individual predation risk by joining group (Reluga, Viscido)
- Public health – Voluntary vaccination games (Bauch, Earn)

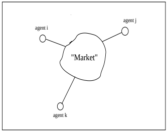

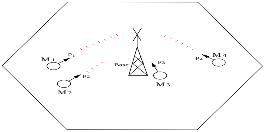

Following the crowd is safer

A generic mean field stochastic control model: N agents

- Dynamics and cost of agent i:

$$dx_i = (1/N)\sum_{j=1}^{N} f_{a_i}(x_i, u_i, x_j)dt + \sigma dw_i, \quad 1 \le i \le N, \quad t \ge 0,$$

$$J_i(u_i, u_{-i}) = E\int_0^T \Big[(1/N)\sum_{j=1}^{N} L(x_i, u_i, x_j)\Big] dt, \quad T < \infty.$$

- Assume independent noises, initial states, etc. u_{-i}: control of other agents. $\{a_i, i \ge 1\}$: depending on individual agents (but having a limiting empirical distribution) to model population heterogeneity
- Other variants of modeling may be considered. In particular, linear quadratic Gaussian (LQG) models allow very explicit analysis

A generic mean field stochastic control model: N agents

An LQG variant –

- Individual dynamics and costs:

$$dx_i = A(\theta_i)x_i dt + Bu_i dt + Cx^{(N)} dt + DdW_i, \qquad 1 \le i \le N,$$
$$J_i = E\int_0^\infty e^{-\rho t}\left\{|x_i - \Phi(x^{(N)})|_Q^2 + u_i^T R u_i\right\} dt,$$

  where $Q \ge 0$, $R > 0$, $\Phi(x^{(N)}) = \Gamma x^{(N)} + \eta$, $|z|_Q^2 = z^T Q z$. J_i may be further modified to use pairwise coupling as in the previous model

- Specification
  - θ_i: dynamic parameter, u_i: control, W_i: noise
  - $x^{(N)} = (1/N)\sum_{i=1}^N x_i$: mean field coupling term

Questions to be asked

If the agents are (I) selfish, or (II) altruistic, what will be a "good" solution framework, respectively?

Potential answers

(I) Look for Nash equilibria (or, furthermore, feedback Nash equilibria by dynamic programming)

(II) Take $J_{\mathrm{soc}}^{(N)} = \sum_{i=1}^{N} J_i$ as an appropriate indicator of group interest. Minimize $J_{\mathrm{soc}}^{(N)}$ as a standard stochastic control problem

However, these approaches are not useful; there is

Recall: $dx_i = (\mathrm{drift})dt + \ldots,$
$J_i = E \int_0^\infty \ldots dt.$

The fundamental difficulties with very large N

By the previous approaches (i.e., standard stochastic dynamic noncooperative games/standard optimal control):

- The computational complexity is too high – the curse of dimensionality
- The informational requirement is demanding – each agent needs to acquire detailed state information of all other agents

This situation suggests the development of a new methodology (using ideas from statistical mechanics)

Related literature

Mean field game models with peers (i.e. comparably small players):

- J.M. Lasry and P.L. Lions (2006a,b, JJM'07): Mean field equilibrium; O. Gueant (JMPA'09)
- G.Y. Weintraub et. el. (NIPS'05, Econometrica'08): Oblivious equilibria for Markov perfect industry dynamics; S. Adlakha, R. Johari, G. Weibtraub, A. Goldsmith (CDC'08): further generalizations with OEs
- M. Huang, P.E. Caines and R.P. Malhame (CDC'03, 04, CIS'06, IEEE TAC'07, SICON'10): Decentralized ε-Nash equilibrium in mean field dynamic games; HCM (Allerton'09, Preprint'11), Asymptotic social optima; M. Nourian, P.E. Caines, et. al. (Preprint'11): mean field consensus model
- T. Li and J.-F. Zhang (IEEE TAC'08): Mean field LQG games with long run average cost

Related literature (ctn)

- H. Yin, P.G. Mehta, S.P. Meyn, U.V. Shanbhag (ACC'10): Games of nonlinear oscillators and phase transition; Mehta et. al. (Preprint'11): application to particle filtering
- H. Tembine et. al. (GameNets'09): Mean field MDP and team; H. Tembine, Q. Zhu, T. Basar (IFAC'11): Risk sensitive mean field games
- V. Kolokoltsov et. al. (SIAM CT'11): fully nonlinear mean field games
- Z. Ma, D. Callaway, I. Hiskens (IEEE CST'12): recharging control of large populations of electric vehicles
- Y. Achdou and I. Capuzzo-Dolcetta (SIAM Numer.'11): Numerical solutions to mean field game equations (coupled PDEs)
- R. Buckdahn, P. Cardaliaguet, M. Quincampoix (DGA'11): Survey
- Rome University Mean Field Game Workshop, May 2011

Related literature (ctn)

Mean field optimal control:

- D. Andersson and B. Djehiche (AMO'11): A new stochastic maximum principle
- J. Yong (Preprint'11): control of mean field Volterra integral equations

Remark: This introduces a different conceptual framework; the control action of the agent has significant impact on the mean field; In contrast, a given agent within a mean field game has little impact on the mean field

Mean field models with a major player:

- H. (SICON'10): LQG models with minor players parameterized by a finite parameter set; develop state augmentation
- S. Nguyen and H. (CDC'11): LQG models with continuum parametrization, Gaussian mean field approximation

von Neumann and Morgenstern (1944, pp. 12)

Very early interest in games with a large number of players –

"... When the number of participants becomes really great, some hope emerges that the influence of every particular participant will become negligible, and that the above difficulties may recede and a more conventional theory becomes possible."

"... In all fairness to the traditional point of view this much ought to be said: It is a well known phenomenon in many branches of the exact and physical sciences that very great numbers are often easier to handle than those of medium size. An almost exact theory of a gas, containing about 10^{25} freely moving particles, is incomparably easier than that of the solar system, made up of 9 major bodies ... This is, of course, due to the excellent possibility of applying the laws of statistics and probability in the first case."

The mean field LQG game

- Individual dynamics:

$$dz_i = (a_i z_i + bu_i)dt + \alpha z^{(N)}dt + \sigma_i dw_i, \quad 1 \le i \le N.$$

- Individual costs:

$$J_i = E\int_0^\infty e^{-\rho t}[(z_i - \Phi(z^{(N)}))^2 + ru_i^2]dt.$$

  - z_i: state of agent i; u_i: control; w_i: noise
    a_i: dynamic parameter; $r > 0$; N: population size
    For simplicity: Take the same control gain b for all agents.
  - $z^{(N)} = (1/N)\sum_{i=1}^N z_i$, Φ: nonlinear function
  - We use this simple scalar model (CDC'03, 04) to illustrate the key idea; generalizations to vector states are obvious

The Nash certainty equivalence (NCE) methodology

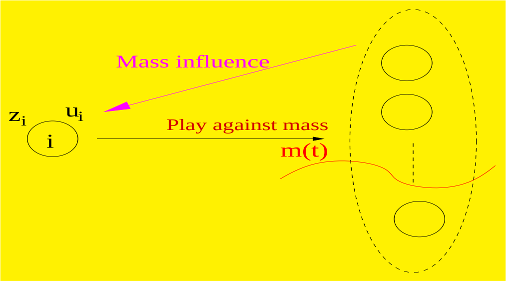

Consistent mean field approximation –

- In the infinite population limit, individual strategies are optimal responses to the mean field $m(t)$;
- Closed-loop behaviour of all agents further replicates the same $m(t)$

The limiting optimal control problem

- Recall

$$dz_i = (a_i z_i + bu_i)dt + \alpha z^{(N)} dt + \sigma_i dw_i$$
$$J_i = E \int_0^\infty e^{-\rho t}[(z_i - \Phi(z^{(N)}))^2 + ru_i^2]dt$$

- Take $f, z^* \in C_b[0,\infty)$ (bounded continuous) and construct

$$d\hat{z}_i = a_i \hat{z}_i dt + bu_i dt + \alpha f dt + \sigma_i dw_i$$
$$J_i(u_i, z^*) = E \int_0^\infty e^{-\rho t}[(\hat{z}_i - z^*)^2 + ru_i^2]dt$$

Riccati Equation : $\rho\Pi_i = 2a_i\Pi_i - (b^2/r)\Pi_i^2 + 1, \qquad \Pi_i > 0.$

- **Optimal Control** : $\widehat{u}_i = -\frac{b}{r}(\Pi_i z_i + s_i)$
  $\rho s_i = \frac{ds_i}{dt} + a_i s_i - \frac{b^2}{r}\Pi_i s_i + \alpha\Pi_i f - z^*.$
- Question: how to determine z^* in reality?

The NCE equation system

Let $\Pi_a = \Pi_i|_{a_i=a}$. Assume (i) $Ez_i(0) = 0$, $i \geq 1$, (ii) The dynamic parameters $\{a_i, i \geq 1\} \subset \mathcal{A}$ have limit empirical distribution $F(a)$.

Optimal control and consistent mean field approximations $\Longrightarrow$

$$\rho s_a = \frac{ds_a}{dt} + as_a - \frac{b^2}{r}\Pi_a s_a + \alpha\Pi_a\bar{z} - z^*,$$

$$\frac{d\bar{z}_a}{dt} = (a - \frac{b^2}{r}\Pi_a)\bar{z}_a - \frac{b^2}{r}s_a + \alpha\bar{z},$$

$$\bar{z} = \int_{\mathcal{A}} \bar{z}_a dF(a),$$

$$z^* = \Phi(\bar{z}).$$

In a system of N agents, agent i uses its own parameter a_i to determine

$$u_i = -\frac{b}{r}(\Pi_{a_i} z_i + s_{a_i}), \quad 1 \leq i \leq N$$

It is decentralized!

Main results: existence, and ε-Nash equilibrium

Theorem (**Existence and Uniqueness**) Under mild assumptions, the NCE equation system has a unique bounded solution $(\bar{z}_a, s_a)$, $a \in \mathcal{A}$.

Let $s_i = s_{a_i}$ be pre-computed from the NCE equation system and

$$u_i^0 = -\frac{b}{r}(\Pi_i z_i + s_i), \quad 1 \le i \le N.$$

Theorem (**Nash equilibria**, CDC'03, TAC'07) The set of strategies $\{u_i^0, 1 \le i \le N\}$ results in an ε-Nash equilibrium w.r.t. costs $J_i(u_i, u_{-i})$, $1 \le i \le N$, i.e.,

$$J_i(u_i^0, u_{-i}^0) - \varepsilon \le \inf_{u_i} J_i(u_i, u_{-i}^0) \le J_i(u_i^0, u_{-i}^0)$$

where $0 < \varepsilon \to 0$ as $N \to \infty$, and u_i depends on $(t, z_1, \ldots, z_N)$.

The nonlinear case

For the nonlinear diffusion model (CIS'06):

- HJB equation:

$$\frac{\partial V}{\partial t} = \inf_{u \in U} \left\{ f[x, u, \mu_t] \frac{\partial V}{\partial x} + L[x, u, \mu_t] \right\} + \frac{\sigma^2}{2} \frac{\partial^2 V}{\partial x^2}$$
$$V(T, x) = 0, \qquad (t, x) \in [0, T) \times \mathbb{R}.$$
$$\Downarrow$$

$$\text{Optimal Control}: \quad u_t = \varphi(t, x|\mu_\cdot), \quad (t, x) \in [0, T] \times \mathbb{R}.$$

- Closed-loop McK-V equation (which can be written as Fokker-Planck equation):

$$dx_t = f[x_t, \varphi(t, x|\mu_\cdot), \mu_t]dt + \sigma dw_t, \quad 0 \le t \le T.$$

The NCE methodology amounts to finding a solution (x_t, μ_t) in McK-V sense.

The model

- Individual dynamics (N agents):

$$dx_i = A(\theta_i)x_i dt + Bu_i dt + DdW_i, \qquad 1 \le i \le N.$$

- Individual costs:

$$J_i = E\int_0^\infty e^{-\rho t}\left\{|x_i - \Phi(x^{(N)})|_Q^2 + u_i^T R u_i\right\} dt,$$

where $\Phi(x^{(N)}) = \Gamma x^{(N)} + \eta$

- Specification
  - θ_i: dynamic parameter, u_i: control, W_i: noise
  - $x^{(N)} = (1/N)\sum_{i=1}^N x_i$: mean field coupling term
- The social cost:

$$J_{\text{soc}}^{(N)} = \sum_{i=1}^N J_i$$

Assumptions

(**A1**) There exists a distribution function $F(\theta)$ on $\mathbb{R}^\kappa$ such that F_N converges to F weakly, i.e., for any bounded and continuous function $\varphi(\theta)$ on $\mathbb{R}^\kappa$,

$$\lim_{N\to\infty} \int \varphi(\theta) dF^{(N)}(\theta) = \int \varphi(\theta) dF(\theta).$$

$\diamondsuit$

(**A2**) The initial states $\{x_i(0), 1 \le i \le N\}$ are independent, $Ex_i(0) = m_0$ for a fixed m_0 and all $i \ge 1$, and there exists $c_0 < \infty$ independent of N such that $\sup_{i\ge1} E|x_i(0)|^2 \le c_0$. $\diamondsuit$

(**A3**) $A(\theta)$ is a continuous matrix function of $\theta \in \Theta$, where Θ is a compact subset of $\mathbb{R}^\kappa$. $\diamondsuit$

(**A4**) For $\theta \in \Theta$, (i) the pair $[A(\theta) - (\rho/2)I, B]$ is stabilizable, (ii) the pair $[Q^{1/2}, A(\theta) - (\rho/2)I]$ is detectable. $\diamondsuit$

The SCE equation system

- The Social Certainty Equivalence (SCE) equation system:

$$\rho s_\theta = \frac{ds_\theta}{dt} + (A_\theta^T - \Pi_\theta BR^{-1}B^T)s_\theta - [(\Gamma^T Q + Q\Gamma - \Gamma^T Q\Gamma)\bar{x} + (I - \Gamma^T)Q\eta],$$

$$\frac{d\bar{x}_\theta}{dt} = A_\theta \bar{x}_\theta - BR^{-1}B^T(\Pi_\theta \bar{x}_\theta + s_\theta),$$

$$\bar{x} = \int \bar{x}_\theta dF(\theta),$$

where $\bar{x}_\theta(0) = m_0$ and s_θ is sought within $C_{\rho/2}([0,\infty), \mathbb{R}^n)$.

- How to construct it? Key idea: combine person-by-person optimality principle in team decision theory with consistent mean field approximations
- With uniform agents (i.e., identical θ) and $\Gamma^T Q + Q\Gamma - \Gamma^T Q\Gamma > 0$, the solution $(s_\theta, \bar{x}_\theta)$ may be obtained in a closed form

The social optimality theorem

(**A5**) There exists a solution $(s_\theta, \bar{x}_\theta, \bar{x}, \theta \in \Theta)$ to the SCE equation system such that each component of $(s_\theta, \bar{x}_\theta, \bar{x})$, as a function of t, is within $C_{\rho/2}([0,\infty), \mathbb{R}^n)$ (i.e., continuous, $O(e^{(\rho-\epsilon)t/2})$ for some $\epsilon > 0$) and such that both s_θ and $\bar{x}_\theta$ are continuous in θ for each fixed $t \in [0,\infty)$. $\diamondsuit$

Theorem Assume (i) (**A1**)-(**A3**), (**A4**)-(i) and (**A5**) hold; (ii) $Q > 0$ and $I - \Gamma$ is nonsingular. Then the set of SCE based control laws

$$\hat{u}_i = -R^{-1}B^T(\Pi_{\theta_i}\hat{x}_i + s_{\theta_i}), \qquad 1 \le i \le N$$

has asymptotic social optimality, i.e., for $\hat{u} = (\hat{u}_1, \dots, \hat{u}_N)$,

$$|(1/N)J^{(N)}_{\mathrm{soc}}(\hat{u}) - \inf_{u\in\mathcal{U}_o}(1/N)J^{(N)}_{\mathrm{soc}}(u)| = O(1/\sqrt{N} + \bar{\epsilon}_N),$$

where $\lim_{N\to\infty}\bar{\epsilon}_N = 0$ and $\mathcal{U}_o$ is defined as a set of centralized information based controls. $\square$

Cost blow-up in the game

The dynamics: $dx_i = ax_i dt + bu_i dt + DdW_i$. The individual cost $J_i = E \int_0^\infty e^{-\rho t}[(x_i - \gamma(x^{(N)} + \eta_0))^2 + ru_i^2]dt, 1 \le i \le N$.
Take $[a; b; D; \rho; r; \eta_0] = [1; 1; 1; 0.8; 0.25; 0.2]$.
We compute the per agent cost when $N \to \infty$.

- γ indicates interaction strength of an individual w. the mean field
- cost blow-up occurs (at $\gamma = 1.09$) in the mean field game, but not in social optimization

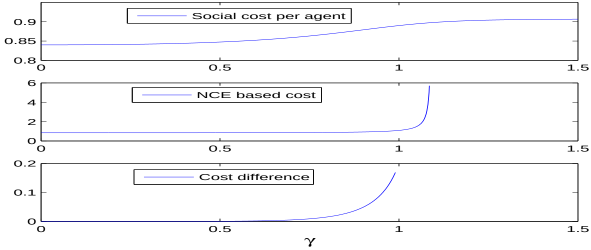

Dynamics with a major player

The LQG game with mean field coupling:

$$dx_0(t) = \left[A_0 x_0(t) + B_0 u_0(t) + F_0 x^{(N)}(t)\right] dt + D_0 dW_0(t), \quad t \geq 0,$$
$$dx_i(t) = \left[A(\theta_i) x_i(t) + B u_i(t) + F x^{(N)}(t) + G x_0(t)\right] dt + D dW_i(t),$$

$x^{(N)} = \frac{1}{N}\sum_{i=1}^N x_i$ mean field term (average state of minor players).

- Major player $\mathcal{A}_0$ with state $x_0(t)$, minor player $\mathcal{A}_i$ with state $x_i(t)$.
- W_0, W_i are independent standard Brownian motions, $1 \leq i \leq N$.
- All constant matrices have compatible dimensions.
- Underlying filtration: $(\Omega, \mathcal{F}, \mathcal{F}_t, t \geq 0, P)$.

We introduce the following assumption:
(**A1**) θ_i takes its value from a finite set $\Theta = \{1, \dots, K\}$ with an empirical distribution $F^{(N)}$, which converges when $N \to \infty$.

Individual costs

The cost for $\mathcal{A}_0$:

$$J_0(u_0, \dots, u_N) = E \int_0^\infty e^{-\rho t} \Big\{ \big| x_0 - \Phi(x^{(N)}) \big|^2_{Q_0} + u_0^T R_0 u_0 \Big\} dt,$$

$\Phi(x^{(N)}) = H_0 x^{(N)} + \eta_0$: cost coupling term

The cost for $\mathcal{A}_i$, $1 \le i \le N$:

$$J_i(u_0, \dots, u_N) = E \int_0^\infty e^{-\rho t} \Big\{ \big| x_i - \Psi(x_0, x^{(N)}) \big|^2_{Q} + u_i^T R u_i \Big\} dt,$$

$\Psi(x_0, x^{(N)}) = H x_0 + \hat{H} x^{(N)} + \eta$: cost coupling term.

- The presence of x_0 in the dynamics and cost of $\mathcal{A}_i$ shows the strong influence of the major player $\mathcal{A}_0$.
- All deterministic constant matrices or vectors H_0, H, $\hat{H}$, $Q_0 \ge 0$, $Q \ge 0$, $R_0 > 0$, $R > 0$, η_0 and η have compatible dimensions.

A matter of "sufficient statistics"

By intuition one might conjecture asymptotic Nash equilibrium strategies of the form:

- $u_0(t)$ for the major player: A function of $(t, x_0(t))$.
- In other words:

  $x_0(t)$ *would be sufficient statistic for* $\mathcal{A}_0$*'s decision;*
  $(x_0(t), x_i(t))$ *would be sufficient statistics for* $\mathcal{A}_i$*'s decision.*

Facts:

- The above conjecture fails!
- Main reason: The major player produces strong influence on the evolution of the mean field (which is now random even in the population limit).

State space augmentation method

Approximate $x^{(N)} = \frac{1}{N}\sum_{i=1}^{N} x_i$ by a process $\bar{z}(t)$. The mean field process is assumed to be governed by ($\bar{x}_0$ is the infinite population version of x_0)

$$d\bar{z}(t) = \overline{A}\bar{z}(t)dt + \overline{G}\bar{x}_0(t)dt + \overline{m}(t)dt,$$

where $\bar{z}(0) = 0$, $\overline{A} \in \mathbb{R}^{nK\times nK}$ and $\overline{G} \in \mathbb{R}^{nK\times n}$ are constant matrices, and $\overline{m}(t)$ is a continuous $\mathbb{R}^{nK}$ function on $[0,\infty)$.

- But so far, none of $\overline{A}$, $\overline{G}$ and $\overline{m}(t)$ is known a priori.
  - The difficulty is overcome by consistent mean field approximations
  - Each agent solves a local limiting control problem; in the end their closed-loop replicates $\overline{A}$, $\overline{G}$ and $\overline{m}(t)$

Asymptotic Nash equilibrium

Algebraic conditions may be obtained for solubility to the above procedure.

The control strategies of $\mathcal{A}_0$ and $\mathcal{A}_i$, $1 \le i \le N$:

$$\hat{u}_0 = -R_0^{-1}\mathbb{B}_0^T[P_0(x_0^T, \bar{z}^T)^T + s_0],$$
$$\hat{u}_i = -R^{-1}\mathbb{B}^T\left[P(x_i^T, x_0^T, \bar{z}^T)^T + s\right], \quad 1 \le i \le N,$$

where (the stochastic ODE of $\bar{z}$ will be driven by x_0)

- $\mathbb{B}_0$, $\mathbb{B}$: determined from coefficients in the original SDE.
- P_0, P: obtained from coupled Algebraic Riccati Equations (ARE).
- s_0, s: obtained from a set of linear ODE.

Theorem (H., SICON'10) Under some technical conditions, the set of decentralized strategies is an ε-Nash equilibrium as $N \to \infty$.

A numerical example (H., 2010)

- The dynamics of the major/minor players are given by

$$dx_0 = 2x_0 dt + u_0 dt + 0.2x^{(N)} dt + dW_0,$$
$$dx_i = 3x_i dt + x_0 dt + u_i dt + 0.3x^{(N)} dt + dW_i, \quad 1 \le i \le N,$$

- The parameters in the costs are: discount factor $\rho = 1$,

$$[Q_0, R_0, H_0, \eta_0] = [1, 1, 0.3, 1.5], \qquad [Q, R, H, \hat{H}, \eta] = [1, 3, 0.4, 0.3, 1].$$

- The dynamics for the mean field: $dz = \overline{A}\overline{z}dt + \overline{G}x_0 dt + \overline{m}dt$
- By an iteration algorithm (for NCE approach), we obtain

$$\overline{A} = -2.06819117030469, \quad \overline{G} = -0.12205345839681.$$

Iterations

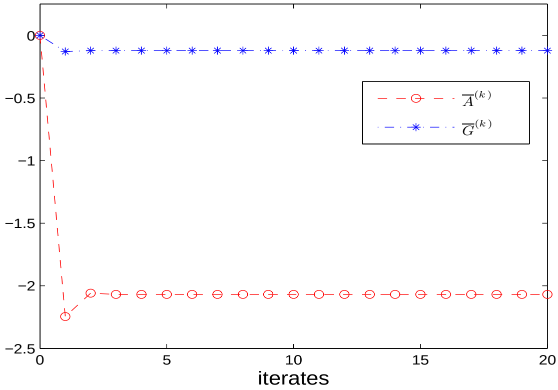

The dynamics for the mean field: $d\bar{z} = \overline{A}\bar{z}dt + \overline{G}x_0dt + \overline{m}dt$

Continuum parametrized minor players

Dynamics and costs:

$$dx_0(t) = \left[A_0 x_0(t) + B_0 u_0(t) + F_0 x^{(N)}(t)\right] dt + D_0 dW_0(t),$$
$$dx_i(t) = \left[A(\theta_i) x_i(t) + B(\theta_i) u_i(t) + F(\theta_i) x^{(N)}(t)\right] dt + D(\theta_i) dW_i(t),$$

$$J_0(u_0, u_{-0}) = E \int_0^T \left[\left|x_0(t) - \chi_0(x^{(N)}(t))\right|^2_{Q_0} + u_0^T(t) R_0 u_0(t)\right] dt,$$

where $\chi_0(x^{(N)}(t)) = H_0 x^{(N)}(t) + \eta_0$, $Q_0 \geq 0$ and $R_0 > 0$.

$$J_i(u_i, u_{-i}) = E \int_0^T \left[\left|x_i(t) - \chi(x_0(t), x^{(N)}(t))\right|^2_{Q} + u_i^T(t) R u_i(t)\right] dt,$$

where $\chi(x_0(t), x^{(N)}(t)) = H x_0(t) + \hat{H} x^{(N)}(t) + \eta$, $Q \geq 0$ and $R > 0$.

Assumptions (with θ being from a continuum set)

(**A1**) The initial states $\{x_j(0), 0 \le j \le N\}$, are independent, and there exists a constant C independent of N such that $\sup_{0\le j\le N} E|x_j(0)|^2 \le C$.
(**A2**) There exists a distribution function $\mathbf{F}(\theta, x)$ on $\mathbb{R}^{d+n}$ such that the sequence of empirical distribution functions $\mathbf{F}_N(\theta, x) = \frac{1}{N}\sum_{i=1}^N 1_{\{\theta_i \le \theta, Ex_i(0)\le x\}}$, $N \ge 1$, where each inequality holds componentwise, converges to $\mathbf{F}(\theta, x)$ weakly, i.e., for any bounded and continuous function $h(\theta, x)$ on $\mathbb{R}^{d+n}$,

$$\lim_{N\to\infty} \int_{\mathbb{R}^{d+n}} h(\theta, x) d\mathbf{F}_N(\theta, x) = \int_{\mathbb{R}^{d+n}} h(\theta, x) d\mathbf{F}(\theta, x).$$

(**A3**) $A(\cdot), B(\cdot), F(\cdot)$ and $D(\cdot)$ are continuous matrix-valued functions of $\theta \in \Theta$, where Θ is a compact subset of $\mathbb{R}^d$.

The Gaussian mean field approximation

- The previous finite dimensional sub-mean field approximations are not applicable
- We consider Gaussian mean field approximations and use a kernel representation:

$$(x^{(N)} \approx) z(t) = f_1(t) + f_2(t) x_0(0) + \int_0^t g(t,s) dW_0(s),$$

where f_1, f_2, g are continuous vector/matrix functions of t

- Individual agents solve limiting control problems with random coefficient processes by a linear BSDE approach
- Consistency condition is imposed for mean field approximations
- The resulting decentralized strategies are not Markovian
- ε-Nash equilibrium can be established (Nguyen and H., CDC'11)

An example

The dynamics and costs:

$$dx_0 = [a_0 x_0(t) + b_0 u_0(t)]dt + D_0 dw_0(t), \quad t \geq 0,$$
$$dx_i = [a_i x_i(t) + b u_i(t)]dt + D dw_i(t), \quad 1 \leq i \leq N,$$

$$J_0(u_0, u_{-0}) = E\int_0^T \left\{q_0(x_0(t) - h_0 x^{(N)}(t) - \eta_0)^2 + u_0^2(t)\right\}dt,$$
$$J_i(u_i, u_{-i}) = E\int_0^T \left\{q(x_i(t) - h x_0(t) - \hat{h} x^{(N)}(t) - \eta)^2 + u_i^2(t)\right\}dt.$$

$[a_0, b_0, D_0, q_0, h_0, \eta_0] = [0.5, 1, 1, 1, 0.6, 1.5]$, $[\underline{a}, \bar{a}, b, D, q, h, \hat{h}, \eta] = [0.1, 0.4, 1, 1, 1.2, 0.5, 0.4, 0.5]$. The empirical distribution of $\{a_i, i \geq 1\}$ converges to uniform distribution on $[\underline{a}, \bar{a}]$.

Numerical solution

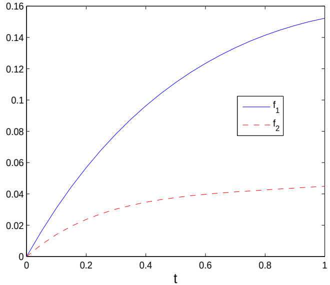

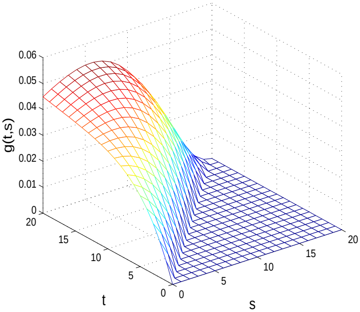

$$z(t) = f_1(t) + f_2(t)x_0(0) + \int_0^t g(t,s)dW_0(s).$$